

A Time Lower Bound for Multiple Nucleation on a Surface

Aaron Sterling*

August 4, 2009

Abstract

Majumder, Reif and Sahu have presented a stochastic model of reversible, error-permitting, two-dimensional tile self-assembly, and showed that restricted classes of tile assembly systems achieved equilibrium in (expected) polynomial time. One open question they asked was how much computational power would be added if the model permitted multiple nucleation, *i.e.*, independent groups of tiles growing before attaching to the original seed assembly. This paper provides a partial answer, by proving that if a tile assembly model uses only local binding rules, then it cannot use multiple nucleation on a surface to solve certain “simple” problems in constant time (time independent of the size of the surface). Moreover, this time bound applies to macroscale robotic systems that assemble in a three-dimensional grid, not just to tile assembly systems on a two-dimensional surface. The proof technique defines a new model of distributed computing that simulates tile (and robotic) self-assembly. *Keywords:* self-assembly, multiple nucleation, locally checkable labeling.

1 Introduction

1.1 Overview

Nature is replete with examples of the self-assembly of individual parts into a more complex whole, such as the development from zygote to fetus, or, more simply, the replication of DNA itself. In his Ph.D. thesis in 1998, Winfree proposed a formal mathematical model to reason algorithmically about processes of self-assembly [21]. Winfree connected the experimental work of Seeman [16] (who had built “DNA tiles,” molecules with unmatched DNA base pairs protruding in four directions, so they could be approximated by squares with different “glues” on each side) to a notion of tiling the integer plane developed by Wang in the

*Laboratory for Nanoscale Self-Assembly, Department of Computer Science, Iowa State University, Ames, IA 50014, USA. sterling@cs.iastate.edu. This research was supported in part by National Science Foundation Grants 0652569 and 0728806.

1960s [20]. Rothmund, in his own Ph.D. thesis, extended Winfree’s original Tile Assembly Model [14].

Informally speaking, Winfree effectivized Wang tiling, by requiring a tiling of the plane to start with an individual *seed tile* or a connected, finite *seed assembly*. Tiles would then accrete one at a time to the seed assembly, growing a *seed supertile*. A *tile assembly system* is a finite set of *tile types*. Tile types are characterized by the names of the “glues” they carry on each of their four sides, and the binding strength each glue can exert. We assume that when the tiles interact “in solution,” there are infinitely many tiles of each tile type. Tile assembly proceeds in discrete stages. At each stage s , from all possibilities of tile attachment at all possible locations (as determined by the glues of the tile types and the binding requirements of the system overall), one tile will bind, with tile type and location “chosen” nondeterministically from possible legal bonds at that stage. (Later, we will generalize this so multiple tiles can bind concurrently, at a given stage.) Winfree proved that his Tile Assembly Model is Turing universal.

The abstract Tile Assembly Model (aTAM) is *error-free* and *irreversible*—tiles always bind correctly, and, once a tile binds, it can never unbind. Adleman *et al.* were the first to define a notion of time complexity for tile assembly, using a one-dimensional error-permitting, reversible model, where tiles would assemble in a line with some error probability, then be scrambled, and fall back to the line [1]. Adleman *et al.* proved bounds on how long it would take such models to achieve equilibrium. Majumder, Reif and Sahu have recently presented a two-dimensional stochastic model for self-assembly [11], and have shown that some tiling problems in their model correspond to *rapidly mixing Markov chains*—Markov chains that reach stationary distribution in time polynomial in the state space of legally reachable assemblies.

While the aTAM is nondeterministic, real-world chemical reactions are probabilistic, and discrete molecular interactions are often modeled stochastically. We will define a class of stochastic self-assembly models that contains the model of Majumder *et al.*, and prove a lower bound about any model in that class.

The tile assembly systems analyzed in [11] had the property that their equilibrium assemblies were identical (allowing for small error) with their *terminal* or *complete assemblies*, *i.e.*, assemblies that cannot legally evolve further, given the rules of the system. This identity does not, however, hold in general. In a closed chemical system, where equilibrium may be achieved, it is possible that the system at equilibrium might consist almost entirely of large, undesirable assemblies that do not perform the desired computation. In these cases, correct assembly occurs when the system is out of equilibrium, and can be maintained because there is a large kinetic energy barrier to forming undesired structures. Therefore, when we discuss the “solution to a problem” in this paper, we identify that with the notion of a complete assembly.

We will prove a time complexity lower bound on the solution of a graph coloring problem for a class of self-assembly models, including, but not limited to, a generalization of the model of [11]. The tile assembly model in [11], like the aTAM, allows only for a single seed assembly, and one of the open problems

in [11] was how the model might change if it allowed multiple nucleation, *i.e.*, if multiple supertiles could build independently before attaching to a growing seed supertile. The main result of this paper provides a time complexity lower bound for a class of tile assembly models that permit multiple nucleation on a 2D surface or a 3D grid: there is no way for those models to use multiple nucleation to achieve a speedup to tiling a surface in constant time (time independent of the size of the surface) in order to solve a graph coloring problem, even though that graph coloring problem requires only seven tile types to solve in the aTAM. This result holds for tile assembly models that are reversible, irreversible, error-permitting or error-free. In fact, a speedup to constant time is impossible, even if we relax the model to allow that, at each step s , there is a positive probability for every available location that a tile will bind there (instead of requiring that exactly one tile bind per stage).

To our knowledge, the method of proof in this paper is novel: given a tile assembly model and a tile assembly system \mathcal{T} in that model, we construct a distributed network of processors that can simulate the behavior of \mathcal{T} as it assembles on a surface. Our result then follows from the theorem by Naor and Stockmeyer that locally checkable labeling (LCL) problems have no local solution in constant time [12]. This is true for both deterministic and randomized algorithms, so no constant-time tile assembly system exists that solves an LCL problem with a positive probability of success. We consider one LCL problem in specific, the weak c -coloring problem, and demonstrate a tile set of only seven tile types that solves the weak c -coloring problem in the abstract Tile Assembly Model, even though that same problem is impossible to solve in constant time by multiple nucleation on a surface, for a broad class of self-assembly models. Intuitively, this demonstrates that even a problem that can be solved in polynomial time by using a few local rules when starting from a single point, cannot necessarily be solved in constant time when starting from multiple points, regardless of the rule set used. (The abstract Tile Assembly Model can weakly c -color an $n \times n$ surface in n^2 steps, yet none of the multiple nucleation models we consider can solve the weak c -coloring problem in constant-many steps.)

The results of Naor and Stockmeyer we apply are more powerful than needed to obtain the time complexity lower bound for a system in which the self-assembling agents are as simple as DNA tiles. Our lower bound actually demonstrates that constant-time speedup to solve LCL problems is impossible via multiple nucleation, even for self-assembling modular robots capable of forming physical bonds in a three-dimensional grid, and, in addition, of sending messages to their neighbors once they have bonded, and potentially deciding to break bonds they previously formed.

1.2 Background

In the abstract Tile Assembly Model, one tile is added per stage, so the primary complexity measure is not one of time, but of how much information a tile set needs in order to solve a particular problem. Several researchers [1] [3] [4] [15] [17] have investigated the *tile complexity* (the minimum number of distinct tile types

required for assembly) of finite shapes, and sets of “scale-equivalent” shapes (essentially a $\mathbb{Z} \times \mathbb{Z}$ analogue of the Euclidean notion of similar figures). For example, it is now known that the number of tile types required to assemble a square of size $n \times n$ (for n any natural number) is $\Omega(\log n / \log \log n)$ [15]. Or, if T is the set of all discrete equilateral triangles, the asymptotically optimal relationship between triangle size and number of tiles required to assemble that triangle, is closely related to the Kolmogorov Complexity of a program that outputs the triangle as a list of coordinates [17].

Despite these advances in understanding of the complexity of assembling finite, bounded shapes, the self-assembly of infinite structures is not as well understood. In particular, there are few lower bounds or impossibility results on what infinite structures can be self-assembled in the Tile Assembly Model. The first such impossibility result appeared in [10], when Lathrop, Lutz and Summers showed that no finite tile set can assemble the discrete Sierpinski Triangle by placing a tile only on the coordinates of the shape itself. (By contrast, Winfree had shown that just seven tile types are required to tile the first quadrant of the integer plane with tiles of one color on the coordinates of the discrete Sierpinski Triangle, and tiles of another color on the coordinates of the complement [21].) Recently, Patitz and Summers have extended this initial impossibility result to other discrete fractals [13], and Lathrop *et al.* [9] have demonstrated sets in $\mathbb{Z} \times \mathbb{Z}$ that are Turing decidable but cannot be self-assembled in Winfree’s sense.

To date, there has been no work comparing the strengths of different tile assembly models with respect to infinite (nor to finite but arbitrarily large) structures. Since self-assembly is a process in which each point has only local knowledge, it is natural to consider whether the techniques of distributed computing might be useful for comparing models of self-assembly and proving impossibility results about them. This paper is an initial attempt in that direction.

Aggarwal *et al.* in [3] proposed a generalization of the standard Tile Assembly Model, which they called the q -Tile Assembly Model. This model permitted multiple nucleation: tiles did not need to bind immediately to the seed supertile. Instead, they could form independent supertiles of size up to some constant q before then attaching to the seed supertile. While the main question considered in [3] was *tile* complexity, we can also ask whether multiple nucleation would allow an improvement in *time* complexity. Intuitively, Does starting from multiple points allow us to build things strictly faster than starting from a single point?

As mentioned above, Majumder, Reif and Sahu recently presented a stochastic, error-permitting tile assembly model, and calculated the rate of convergence to equilibrium for several tile assembly systems [11]. The model in [11] permitted only a single seed assembly, and addition of one tile to the seed supertile at each stage. Majumder, Reif and Sahu left as an open question how the model might be extended to permit the presence and binding of multiple supertiles.

Therefore, we can rephrase the “intuitive” question above as follows: Can we tile a surface of size $n \times n$ in a constant number of stages, by randomly selecting nucleation points on the surface, building supertiles of size q or smaller from

those points in $\leq q$ stages, and then allowing $\leq r$ additional stages for tiles to fall off and be replaced if the edges of the supertiles contain tiles that bind incorrectly? (The assembly achieves equilibrium in constant time because q and r do not depend on n .) The partial answer obtained in this paper is that locally checkable labeling problems cannot be solved in constant time, if we limit ourselves to self-assembly on a surface.

Limiting ourselves to self-assembly on a surface is significant, because we are requiring that agents adhere to a substrate and then never move again, unless they dissociate completely from the larger assembly. When assemblies multiply nucleate in solution, however, they form disjoint supertiles that can float independently until potentially becoming aligned, with some probability. A self-assembly model that made this rigorous might be strictly stronger than the self-assembly models we consider in this paper, as it is not clear how to simulate floating supertiles within our distributed computing models without introducing slowdown, as processors simulating locations of the surface would have to “pass along” information from one processor to the next, to simulate elements of the moving supertile. We leave the possibility of simulating floating supertiles to future work.

Another limitation to our results is that our proof technique applies only to self-assembly models whose binding rules are completely local. One could imagine models in which supertiles combine (or separate) based on simultaneous interactions at several locations, instead of the models we consider in this paper, in which the system’s behavior at each location depends only on the properties of that location’s immediate neighbors. The self-assembly literature, to our knowledge, contains little regarding self-assembly models with nonlocal binding rules, and this could be a fruitful area to investigate.

Klavins and co-authors have modeled self-assembly phenomena—and programmed self-assembling modular robots—using graph grammars [6] [8]. Klavins in [7] informally compares the limitations of the “distributed algorithms” of graph grammars (used to program self-assembling robots) to impossibility results in distributed computing. Recently, we have shown connections between self-assembly and the wait-free consensus hierarchy [18], and we have embedded the “graph assembly systems” of Klavins into a known graph grammar characterization of distributed systems [19]. The present paper, to the best of our knowledge, is the first to construct a formal reduction from self-assembly models to models of distributed computing.

Section 2 of this paper describes the abstract Tile Assembly Model, and then considers generalizations of the standard model that permit multiple nucleation. Section 3 reviews the distributed computing results of Naor and Stockmeyer needed to prove the impossibility result. In Section 4 we present our simulation technique and lower bound results. Section 5 concludes the paper and suggests directions for future research.




Figure 1: An example tile with explanation.

2 Description of Self-Assembly Models

2.1 The abstract Tile Assembly Model

Winfree’s objective in defining the Tile Assembly Model was to provide a useful mathematical abstraction of DNA tiles combining in solution [21]. Rothemund [14], and Rothemund and Winfree [15], extended the original definition of the model. For a comprehensive introduction to tile assembly, we refer the reader to [14]. In our presentation here, we follow [10], which gives equal status to finite and infinite tile assemblies.

Intuitively, a tile of type t is a unit square that can be placed with its center on a point in the integer lattice. A tile has a unique orientation; it can be translated, but not rotated. We identify the side of a tile with the direction (or unit vector) one must travel from the center to cross that side. The literature often refers to west, north, east and south sides, starting at the leftmost side and proceeding clockwise. Each side of a tile is covered with a “glue” that has a *color* and a *strength*. Figure 1 shows how a tile is represented graphically.

If tiles of types t and t' are placed adjacent to each other, then they will *bind* with the strength shared by both adjacent sides if the glues on those sides are the same. Note that this definition of binding implies that if the glues of the adjacent sides do not have the same color or strength, then their binding strength is 0. Later, we will permit pairs of glues to have negative binding strength, to model error occurrence and correction.

One parameter in a tile assembly model is the minimum binding strength

required for tiles to bind “stably.” This parameter is usually termed *temperature* and denoted by τ , where $\tau \in \mathbb{N}$.

As we consider only two-dimensional tile assemblies, we limit ourselves to working in $\mathbb{Z}^2 = \mathbb{Z} \times \mathbb{Z}$. U_2 is the set of all unit vectors in \mathbb{Z}^2 .

A *binding function* on an (undirected) graph $G = (V, E)$ is a function $\beta : E \rightarrow \mathbb{N}$. If β is a binding function on a graph $G = (V, E)$ and $C = (C_0, C_1)$ is a cut of G , then the *binding strength* of β on C is

$$\beta_C = \{\beta(e) \mid e \in E, e \in C_0, \text{ and } e \in C_1\}.$$

The *binding strength* of β on G is then $\beta(G) = \min\{\beta_C \mid C \text{ is a cut of } G\}$. Intuitively, the binding function captures the strength with which any two neighbors are bound together, and the binding strength of the graph is the minimum strength of bonds that would have to be severed in order to separate the graph into two pieces.

A *binding graph* is an ordered triple $G = (V, E, \beta)$ where (V, E) is a graph and β is a binding function on (V, E) . If $\tau \in \mathbb{N}$, a binding graph $G = (V, E, \beta)$ is τ -*stable* if $\beta((V, E)) \geq \tau$.

Recall that a *grid graph* is a graph $G = (V, E)$ where $V \subseteq \mathbb{Z} \times \mathbb{Z}$ and every edge $\{\vec{m}, \vec{n}\} \in E$ has the property that $\vec{m} - \vec{n} \in U_2$. We write $[V]^2$ for the set $\{\{v_1, v_2\} \mid v_1 \in V \text{ and } v_2 \in V\}$, i.e., the two-element subsets of V .

Definition 1. A tile type over a (finite) alphabet Σ is a function $t : U_2 \rightarrow \Sigma^* \times \mathbb{N}$. We write $t = (\text{col}_t, \text{str}_t)$, where $\text{col}_t : U_2 \rightarrow \Sigma^*$, and $\text{str}_t : U_2 \rightarrow \mathbb{N}$ are defined by $t(\vec{u}) = (\text{col}_t(\vec{u}), \text{str}_t(\vec{u}))$ for all $\vec{u} \in U_2$.

Definition 2. If T is a set of tile types, a T -configuration is a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$.

Definition 3. The binding graph of a T -configuration $\alpha : \mathbb{Z}^2 \dashrightarrow T$ is the binding graph $G_\alpha = (V, E, \beta)$, where (V, E) is the grid graph given by

$$V = \text{dom}(\alpha),$$

$$E = \{\{\vec{m}, \vec{n}\} \in [V]^2 \mid \vec{m} - \vec{n} \in U_2, \text{col}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) = \text{col}_{\alpha(\vec{n})}(\vec{m} - \vec{n}), \text{ and } \text{str}_{\alpha(\vec{m})}(\vec{n} - \vec{m}) > 0\},$$

and the binding function $\beta : E \rightarrow \mathbb{Z}^+$ is given by $\beta(\{\vec{m}, \vec{n}\}) = \text{str}_{\alpha(\vec{m})}(\vec{n} - \vec{m})$ for all $\{\vec{m}, \vec{n}\} \in E$.

Definition 4. For T a set of tile types, a T -configuration α is *stable* if its binding graph G_α is τ -stable. A τ - T -assembly is a T -configuration that is τ -stable. We write \mathcal{A}_T^τ for the set of all τ - T -assemblies.

Definition 5. Let α and α' be T -configurations.

1. α is a subconfiguration of α' , and we write $\alpha \sqsubseteq \alpha'$, if $\text{dom}(\alpha) \subseteq \text{dom}(\alpha')$ and, for all $\vec{m} \in \text{dom}(\alpha)$, $\alpha(\vec{m}) = \alpha'(\vec{m})$.

2. α' is a single-tile extension of α if $\alpha \sqsubseteq \alpha'$ and $\text{dom}(\alpha') \setminus \text{dom}(\alpha)$ is a singleton set. In this case, we write $\alpha' = \alpha + (\vec{m} \mapsto t)$, where $\{\vec{m}\} = \text{dom}(\alpha') \setminus \text{dom}(\alpha)$ and $t = \alpha'(\vec{m})$.
3. The notation $\alpha \xrightarrow[\tau, T]{1} \alpha'$ means that $\alpha, \alpha' \in \mathcal{A}_T^\tau$ and α' is a single-tile extension of α . (The “1” above the arrow is to denote that a single tile is added at this step.)

Definition 6. Let $\alpha \in \mathcal{A}_T^\tau$.

1. For each $t \in T$, the τ - t -frontier of α is the set

$$\partial_t^\tau \alpha = \left\{ \vec{m} \in \mathbb{Z}^2 \setminus \text{dom}(\alpha) \mid \sum_{\vec{u} \in U_2} \text{str}_t(\vec{u}) \cdot \llbracket \alpha(\vec{m} + \vec{u})(-\vec{u}) = t(\vec{u}) \rrbracket \geq \tau \right\}.$$

2. The τ -frontier of α is the set

$$\partial^\tau \alpha = \bigcup_{t \in T} \partial_t^\tau \alpha.$$

Definition 7. A τ - T -assembly sequence is a sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ in \mathcal{A}_T^τ , where $k \in \mathbb{Z}^+ \cup \{\infty\}$ and, for each i with $1 \leq i+1 < k$, $\alpha_i \xrightarrow[\tau, T]{1} \alpha_{i+1}$.

Definition 8. The result of a τ - T -assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ is the unique T -configuration $\alpha = \text{res}(\vec{\alpha})$ satisfying: $\text{dom}(\alpha) = \bigcup_{0 \leq i < k} \text{dom}(\alpha_i)$ and $\alpha_i \sqsubseteq \alpha$ for each $0 \leq i < k$.

Definition 9. Let $\alpha, \alpha' \in \mathcal{A}_T^\tau$. A τ - T -assembly sequence from α to α' is a τ - T -assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i < k)$ such that $\alpha_0 = \alpha$ and $\text{res}(\vec{\alpha}) = \alpha'$. We write $\alpha \xrightarrow[\tau, T]{} \alpha'$ to indicate that there exists a τ - T -assembly from α to α' .

Definition 10. An assembly $\alpha \in \mathcal{A}_T^\tau$ is terminal if $\partial^\tau \alpha = \emptyset$.

Intuitively, a configuration is a set of tiles that have been placed in the plane, and the configuration is stable if the binding strength at every possible cut is at least as high as the temperature of the system. Informally, an assembly sequence is a sequence of single-tile additions to the frontier of the assembly constructed at the previous stage. Assembly sequences can be finite or infinite in length. We are now ready to present a definition of a tile assembly system.

Definition 11. Write \mathcal{A}_T^τ for the set of configurations, stable at temperature τ , of tiles whose tile types are in T . A tile assembly system is an ordered triple $\mathcal{T} = (T, \sigma, \tau)$ where T is a finite set of tile types, $\sigma \in \mathcal{A}_T^\tau$ is the seed assembly, and $\tau \in \mathbb{N}$ is the temperature. We require $\text{dom}(\sigma)$ to be finite.

Definition 12. Let $\mathcal{T} = (T, \sigma, \tau)$ be a tile assembly system.

1. Then the set of assemblies produced by \mathcal{T} is

$$\mathcal{A}[\mathcal{T}] = \{ \alpha \in \mathcal{A}_T^\tau \mid \sigma \xrightarrow{\tau, T} \alpha \} ,$$

where “ $\sigma \xrightarrow{\tau, T} \alpha$ ” means that tile configuration α can be obtained from seed assembly σ by a legal addition of tiles.

2. The set of terminal assemblies produced by \mathcal{T} is

$$\mathcal{A}_\square[\mathcal{T}] = \{ \alpha \in \mathcal{A}[\mathcal{T}] \mid \alpha \text{ is terminal} \} ,$$

where “terminal” describes a configuration to which no tiles can be legally added.

If we view tile assembly as the programming of matter, the following analogy is useful: the seed assembly is the input to the computation; the addition of tile types to the growing assembly are the legal steps the computation can take; the temperature is the primary inference rule of the system; and the terminal assemblies are the possible outputs.

We are, of course, interested in being able to *prove* that a certain tile assembly system always achieves a certain output. In [17], Soloveichik and Winfree presented a strong technique for this: local determinism.

Informally, an assembly sequence $\vec{\alpha}$ is locally deterministic if (1) each tile added in $\vec{\alpha}$ binds with the minimum strength required for binding; (2) if there is a tile of type t_0 at location \vec{m} in the result of α , and t_0 and the immediate “OUT-neighbors” of t_0 are deleted from the result of α , then no other tile type in \mathcal{T} can legally bind at \vec{m} ; the result of α is terminal. We formalize these points as follows.

Definition 13 (Soloveichik and Winfree [17]). A τ - \mathcal{T} -assembly sequence $\vec{\alpha} = (\alpha_i \mid 0 \leq i \leq k)$ with result α is locally deterministic if it has the following three properties.

1. For all $\vec{m} \in \text{dom}(\alpha) - \text{dom}(\alpha_0)$,

$$\sum_{\vec{u} \in \text{IN}^{\vec{\alpha}}(\vec{m})} \text{str}_{\alpha_{i_\alpha(\vec{m})}}(\vec{m}, \vec{u}) = \tau ,$$

where $\text{IN}^{\vec{\alpha}}(\vec{m})$ means the sides of the tile that bound at location \vec{m} during assembly sequence $\vec{\alpha}$ that contributed nonzero strength during the stage at which the tile bound. (Informally, these are the “input sides” of the tile at location \vec{m} , with respect to assembly sequence $\vec{\alpha}$.)

2. For all $\vec{m} \in \text{dom}(\alpha) - \text{dom}(\alpha_0)$ and all $t \in \mathcal{T} - \{\alpha(\vec{m})\}$, $\vec{m} \notin \partial_t^r(\vec{\alpha} \setminus \vec{m})$.

3. $\partial^\tau \alpha = \emptyset$.

Definition 14 (Soloveichik and Winfree [17]). A tile assembly system \mathcal{T} is locally deterministic if there exists a locally deterministic τ - \mathcal{T} -assembly sequence $\alpha = (\alpha_i \mid 0 \leq i < k)$ with $\alpha_0 = \sigma$.

Local determinism is important because of the following result.

Theorem 1 (Soloveichik and Winfree [17]). *If \mathcal{T} is locally deterministic, then \mathcal{T} has a unique terminal assembly.*

2.2 More general self-assembly models

We move now from DNA tiles self-assembling on a two-dimensional surface, to a more general setting, where self-assembling “agents” with the ability not just to bind but also to communicate after binding and potentially unbind, can assemble either in the plane or in three-space. One could think of these agents as (nano- or macroscale) robots that interlock physically, and, after interlocking, can send their neighbors electronic messages of low complexity. Based on receipt of messages, the robots can then decide to break bonds to one or more of their neighbors. Such modular robots have already been implemented in laboratory experiments [7]. Further, these robots may be constructed so each has (at least with high probability) a unique identification code—permitting transmission of strictly more information than is possible in the setting of tile self-assembly, in which tiles do not have unique identifiers.

We will consider generalizations of the abstract Tile Assembly Model that include the following: (1) multiple nucleation; (2) assembly in which glues bind incorrectly according to some error probability; and (3) negative glue strengths, allowing incorrectly bound tiles to be released from the assembly so it is possible for a correctly-binding tile to attach in that location; (4) a third spatial dimension; and (5) tiles can now be “agents,” *i.e.*, finite state machines with algorithms and unique identifiers. We formalize this as follows.

Definition 15. *A d -regular self-assembling agent type T is a finite state machine of form $T = \langle A, (g_1, \dots, g_d) \rangle$, where A is an (deterministic or probabilistic) algorithm and the g_i ’s are finite strings over a finite alphabet (codes for the glue types associated with T) that are hardcoded into the machine. The algorithm A can be null (in the case of passive self-assembly like DNA tiles), or can decide whether to transmit messages of length bounded by a constant to neighbors based on the agent’s interaction with neighboring glue types.*

We will assume that all agent types have identical geometric structure, and their d glues all have the same orientation. For example, in the aTAM, all agent types are unit squares, oriented north, east, south, west. Also, for simplicity of the proof, we will assume that our agents are memoryless. However, because of the generality of the results of Naor and Stockmeyer, our lower bound results would still hold if agents could make active self-assembly decisions based on a history of messages received from neighbors, not just the current messages and glue types of their neighbors.

Definition 16. *The binary relation R is a set of binding rules for the (finite) set of agent types $\{T_i\}_i$ if, for any $(x, y) \in R$, both x and y are glue types that appear in elements of $\{T_i\}_i$.*

Definition 17. *The function β is an assignment of binding strengths for the set of binding rules R , if the domain of β is R , and the range of β is the set of nonnegative integers.*

Definition 18. *\mathcal{M} is a model of d -regular self-assembling agents if $\mathcal{M} = \langle \{T_i\}_i, R, \beta, \tau, \sigma \rangle$, where $\{T_i\}_i$ is a (finite) set of d -regular self-assembling agent types, R is a set of binding rules for $\{T_i\}_i$, β is an assignment of binding strengths for R , τ is the temperature of the system (the threshold binding strength for bonds to be stable), and σ is an initial (finite) seed assembly.*

The algorithm of each agent type may include a variable MY-ID, and we allow the possibility that each agent in the system does, in fact, have a unique identification number. This might be appropriate when modeling robotic self-assembly. In the case of molecular self-assembly, each agent is anonymous. Assembly systems in both the aTAM and the stochastic model of Majumder *et al.* can be defined in this formalism, by giving each agent type an algorithm that performs no instructions, and defining (respectively deterministic or probabilistic) binding relations in a natural way.

To conclude this section, we formalize what it means for a self-assembly model to allow multiple nucleation on a surface.

Definition 19. *Let \mathcal{M} be a model of d -regular self-assembling agents. We say \mathcal{M} allows multiple nucleation if, in addition to the placement of the seed assembly at the initial stage of assembly, there is some probability π_ν that (at the first stage of assembly only) an agent is placed on each location of the surface with probability π_ν . Further, if an agent is placed at location \vec{m} because of multiple nucleation, its agent type is chosen uniformly at random from the space of possible agent types.*

We could allow multiple nucleation to occur at multiple stages during the assembly, not just the first. Again, because of the generality of Naor and Stockmeyer’s results, that would not affect our lower bound proof.

3 Distributed Computing Results of Naor and Stockmeyer

In a well known distributed computing paper, Naor and Stockmeyer investigated whether “locally checkable labeling” problems could be solved over a network of processors in an entirely local manner, where a local solution means a solution arrived at “within time (or distance) independent of the size of the network” [12]. One locally checkable labeling problem Naor and Stockmeyer considered was the *weak c -coloring problem*.

Definition 20 (Naor and Stockmeyer [12]). *For $c \in \mathbb{N}$, a weak c -coloring of a graph is an assignment of numbers from $\{1, \dots, c\}$ (the possible “colors”) to the vertices of the graph such that for every non-isolated vertex v there is at least one neighbor w such that v and w receive different colors. Given a graph G , the weak c -coloring problem for G is to weak c -color the nodes of G .*

In the context of tiling, to solve the weak c -coloring problem for an $n \times n$ surface means tiling the surface so each tile has at least one neighbor (to the north, south, east or west) of a different color. In the next section, we will present a simple solution to the weak c -coloring problem in the abstract Tile Assembly Model. By contrast, Naor and Stockmeyer showed that no local, constant-time algorithm can solve the weak c -coloring problem for grid graphs, nor for k -dimensional meshes, a generalization of grid graphs which we now define.

Definition 21. *A k -dimensional mesh is a graph with vertex set $\{0, 1, \dots, m\}^k$ for some m , such that two vertices are connected by an edge if the \mathcal{L}_1 -distance between them is 1.*

Theorem 2 (Naor and Stockmeyer [12]). *For any natural numbers c , k and t , there is no local algorithm with time bound t that solves the weak c -coloring problem for the class of k -dimensional meshes. (This remains true even if the processors have unique identifiers and can transmit them as part of the local algorithm.)*

A second theorem from the same paper says that randomization does not help. The original result is stronger than the formulation here.

Theorem 3 (Naor and Stockmeyer [12]). *Fix a class \mathcal{G} of graphs closed under disjoint union. If there is a randomized local algorithm P with time bound t that solves the weak c -coloring problem for \mathcal{G} with error probability ϵ for some $\epsilon < 1$, then there is a deterministic local algorithm A with time bound t that solves the weak c -coloring problem for \mathcal{G} .*

4 Simulation of Self-Assembly on a Surface

In order to apply the theorems of Naor and Stockmeyer to the realm of self-assembly, we build a distributed network of processors that reduces a self-assembly problem to a distributed computing problem. The motivating intuition is that each processor simulates a location of the surface, and reports to its neighbors whether there is a (simulated) agent at that location. Formally, we prove the following theorem.

Theorem 4. *Let \mathcal{M} be a model of d -regular self-assembling agents for any natural number $d > 0$, such that \mathcal{M} self-assembles on a k -dimensional mesh of size n^k , and that \mathcal{M} allows multiple nucleation. Then there is a model \mathcal{N} of distributed computing that simulates \mathcal{M} using n^k processors with the network topology of a k -dimensional mesh, and constant-size message complexity.*

Proof. Fix a model of d -regular self-assembling agents \mathcal{M} as in the theorem statement. Let α be a configuration of agents on the mesh of size n^k . Let Γ be the set of glue types of \mathcal{M} and M the set of electronic messages of \mathcal{M} . (Both

Γ and M are finite sets.) The definition of the binding function β induces a function

$$\hat{\beta} : (T \cup \{\emptyset\}) \times (\Gamma \cup \{\emptyset\})^d \times (M \cup \{\emptyset\})^d \times (T \cup \{\emptyset\}) \longrightarrow [0, 1]$$

such that $\hat{\beta}$ takes as input the (possibly empty) agent type $\alpha(\vec{m})$ at some location \vec{m} in configuration α , and, based on the glue types and electronic messages received from the d neighbors that could be incident to an agent at \vec{m} , returns, for each agent type, the probability that $\alpha(\vec{m})$ would contain that agent type, over the space of all legal \mathcal{M} -assembly sequences that start with configuration α and run for one time step. In particular, for fixed $t_0 \in T$, fixed $\gamma_1, \dots, \gamma_d \in \Gamma \cup \{\emptyset\}$, and fixed $m_1, \dots, m_d \in M \cup \{\emptyset\}$, it is true that

$$\sum_{t \in T \cup \{\emptyset\}} \hat{\beta}(t_0, \gamma_1, \dots, \gamma_d, m_1, \dots, m_d, t) = 1 \ .$$

We have not formally defined \mathcal{M} -assembly sequences, but they are a natural extension of the τ - T -assembly sequences of tile self-assembly, where the β and τ of \mathcal{M} are used to determine whether agents bind stably to one another. Also, if an agent type lies on the edge of the n^k -size surface, so it does not have a neighbor in a particular direction, we define $\hat{\beta}$ so that the empty set is the glue and electronic message “transmitted” from the “neighbor” in that direction.

We simulate assembly sequences of \mathcal{M} on an k -dimensional mesh where each of the dimensions has length n by a network of processors \mathcal{N} whose network graph is also a k -dimensional mesh of total size n^k . Each processor will simulate the presence or absence of an agent in the same location on the assembly surface. We interpret bonds between two agents as messages. We add on top of those messages, an additional set of electronic messages agents can send neighbors, and encode the combination as an ordered pair: glue type and electronic message. The function $\hat{\beta}$ will be the probabilistic transition function for processors in this system.

Processors of \mathcal{N} are of the following form.

Processor p_i

d -many input message buffers: $\text{inbuf}_{i,1}, \dots, \text{inbuf}_{i,d}$.

d -many output message buffers: $\text{outbuf}_{i,1}, \dots, \text{outbuf}_{i,d}$.

A color variable: COLOR_i , a variable that can take a value from $\{1, \dots, c\}$, where c is a global constant.

A local state: Each processor is in one of $|T| + 1$ different local states q during a given execution stage s . There is one stage q_k to simulate each agent type $t_k \in T$, and an additional stage EMPTY , to simulate the absence of an agent from the surface location that p_i is simulating.

A state transition function: This function takes the current processor state and the messages received in the current round, and probabilistically directs what state the processor will adopt in the next round.

The messages processors send on the network are of form (glue type, electronic message). The input message buffers of processor p_i simulate the glue types of the edges the agent at p_i 's location is adjacent to, and the electronic messages (if any) received from an agent's neighbors. The output message buffers of p_i simulate the glues on the edges of the tile p_i is simulating, and the electronic messages the agent transmits to its neighbors. The purpose of COLOR_i is to simulate the color of the agent placed at the location simulated by p_i .

All processors in \mathcal{N} are hardcoded with the same probabilistic state transition function, which is determined from the definition of $\hat{\beta}$ (which we induced above from the properties of \mathcal{M}), in the natural way: if, in round r of the algorithm execution, p_i is in state q_k , a simulation of $t_k \in T$, and hears messages that simulate glue types g_1, \dots, g_d and electronic messages m_1, \dots, m_d , then at the end of round r , it will transition to state q_j with probability π_j , where $\hat{\beta}(t_k, g_1, \dots, g_d, m_1, \dots, m_d, t_j) = \pi_j$ and each t_j is a distinct element of $T \cup \{\emptyset\}$. As explained above, we denote the state that simulates the "presence of the empty set"—*i.e.*, the absence of any agent from the location simulated by p_i —as EMPTY .

To simulate the process of self-assembly, we run the following distributed algorithm on \mathcal{N} .

Algorithm execution proceeds in synchronized rounds. Before execution begins, all processors start in state EMPTY . In round $r = 0$, (through the intervention of an omniscient operator) each processor in the locations corresponding to the seed assembly enters the stage to simulate the agent type at that location in the seed assembly.

Also in round $r = 0$, each processor not simulating part of the seed assembly "wakes up" (enters a state other than EMPTY) with probability π_ν , the multiple nucleation probability of \mathcal{M} . If a processor wakes up, it enters state $q \neq \text{EMPTY}$, chosen uniformly at random from the set of non- EMPTY states. For any round $r > 0$, each processor runs either Algorithm 1 or Algorithm 2, depending on whether it is in state EMPTY .

The interaction between agents in \mathcal{M} is completely defined by the glues and electronic messages of an agent's immediate neighbors, as specified in the function $\hat{\beta}$ and the algorithm of each agent type. The processors of \mathcal{N} simulate that behavior with Algorithm 2. Since the processors of \mathcal{N} simulate empty locations with Algorithm 1, by a straightforward induction argument, \mathcal{N} can simulate all possible \mathcal{M} -assembly sequences, and the theorem is proved. \square

We obtain our time lower bound results as corollaries of Theorem 4.

Corollary 1. *If the (deterministic or probabilistic) binding rules of a multiply nucleating tile assembly system \mathcal{T} are entirely local, then \mathcal{T} is unable to solve the weak c -coloring problem in constant time.*

Proof. Suppose \mathcal{T} is an irreversible tiling model. If \mathcal{T} can weak c -color surfaces in constant time, then there is a deterministic algorithm for the distributed

Algorithm 1 For p_i in state EMPTY at round r

```

if  $r = 0$  then
    wake up with probability  $\pi_\nu$ , and cease execution for this round.
end if
if  $r > 0$  then
    Read the  $d$ -many input buffers.
    if no messages were received then
        cease execution for the round
    else
        let  $q_0$  be the state change obtained according to probabilities  $\hat{\beta}$  assigns
        to the space  $T \cup \{\emptyset\}$ , for a location that has adjacent glue types and elec-
        tronic messages that are simulated by the messages received this round.

        Send the messages indicated by state  $q_0$  and the behavior of  $A$ .
        Set the value of  $\text{COLOR}_i$  according to  $q_0$ .
        Enter state  $q_0$  and cease execution for this round.
    end if
end if

```

Algorithm 2 For p_i in state $q \neq \text{EMPTY}$ (at any round)

```

    Read the four input buffers.
    if no messages were received then
        Send the messages indicated by state  $q$  and the behavior of  $A$  and cease
        execution for this round.
    else
        Let  $q_0$  be the state change obtained probabilistically, based on the proba-
        bilities produced by the function  $\hat{\beta}$  to the space  $T \cup \{\emptyset\}$ , given input from
        the glue types and electronic messages simulated by the messages received
        this round.
        Send the messages indicated by state  $q_0$ .
        Set the value of  $\text{COLOR}_i$  according to  $q_0$ .
        Enter state  $q_0$  and cease execution for this round.
    end if

```




Figure 2: The tileset \mathcal{T}^* used in the proof of Proposition 1.

network \mathcal{N} that weak c -colors \mathcal{N} locally, and in constant time. By Theorem 2 that is impossible.

So assume \mathcal{T} is a reversible tiling model, and when \mathcal{T} assembles, it weak c -colors the tiling surface, and achieves bond pair equilibrium in constant time. Then there is a local probabilistic algorithm for \mathcal{N} that weak c -colors \mathcal{N} in constant time, with positive probability of success. By Theorem 3 that is impossible as well. Therefore, no \mathcal{T} exists that weak c -colors surfaces in constant time. \square

By a similar argument, we obtain a lower bound for active self-assembling agents on a three-dimensional cubic grid.

Corollary 2. *If a model of 6-regular self-assembling agents has only local binding rules, then it cannot solve the weak c -coloring problem in constant time on a 3-dimensional mesh, for any value of c .*

A physical interpretation of Corollary 2 would be that robots self-assembling in three-space (*i.e.*, $k = 3$ and, in this example, $d = 6$ so there are six arms coming off each robot, orthogonally to one another) cannot achieve speedup to constant time by self-assembling in separate groups and then joining the groups together. This lower bound remains in effect even if the robots are designed by a method that assigns each robot a unique identifier.

We conclude this section by noting that the weak c -coloring problem has low tile complexity—that is, can be defined using only a few local rules—in the aTAM.

Proposition 1. *There is a tile assembly system in the abstract Tile Assembly Model that weak c -colors the first quadrant, using only seven distinct tile types.*

Proof. Figure 2 exhibits a tiling \mathcal{T}^* of seven tile types that assembles into a weak c -coloring of the first quadrant, starting from an individual seed tile placed at the origin. One can verify by inspection that \mathcal{T}^* is locally deterministic, so it will always produce the same terminal assembly. All assembly sequences generated by \mathcal{T}^* produce a checkerboard pattern in which a monochromatic “+” configuration never appears. Hence, it solves the weak c -coloring problem for the entire first quadrant, and also for all $n \times n$ squares, for any n . \square

One can define a three-dimensional version of the tiling \mathcal{T}^* (shown in Figure 2) in the natural way, using for example the 3D tile assembly model in [5]. Such a three-dimensional tiling will weakly c -color the three-dimensional mesh where $d = 6$, with low tile complexity.

5 Conclusion

In this paper, we showed that if a tile assembly model has only local binding rules, then it cannot use multiple nucleation on a surface to solve locally checkable labeling problems in constant time, even though the abstract Tile Assembly Model can solve a locally checkable labeling problem using just seven tile types. In fact, we proved a more general impossibility result, which showed the same lower bound applies to self-assembling agents in a three-dimensional grid that are capable of binding and subsequently sending messages to their neighbors. To the best of our knowledge, this was the first application of a distributed computing impossibility result to the field of self-assembly.

There are still many open questions regarding multiple nucleation. Aggarwal *et al.* asked in [3] whether multiple nucleation might reduce the tile complexity of finite shapes. The answer is not known. Furthermore, we can ask for what class of computational problems does there exist some function f such that we could tile an $n \times n$ square in time $\mathcal{O}(1) < \mathcal{O}(f) < \mathcal{O}(n^2)$, and “solve” the problem with “acceptable” probability of error, in a tile assembly model that permits multiple nucleation. It would also be interesting to explore the possibility of modeling multiple nucleation of molecules floating in solution—instead of adhering to a surface—perhaps by using techniques from the field of *ad hoc* wireless networks.

We hope that this is just the start of a conversation between researchers in the fields of distributed computing and biomolecular computation.

Acknowledgements

I am grateful to Soma Chaudhuri, Dave Doty, Jim Lathrop and Jack Lutz for helpful discussions on earlier versions of this paper. I am also grateful to two anonymous referees, who suggested significant conceptual and technical improvements to my original journal submission.

References

- [1] Adleman, L., Cheng, Q., Goel, A., Huang, M.D.: Running time and program-size for self-assembled squares. In: Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing, pp. 740-748 (2001)
- [2] Attiya, H., Welch, J.: Distributed Computing: Fundamentals, Simulations, and Advanced Topics, second edition. Wiley Series on Parallel and Distributed Computing (2004)
- [3] Aggarwal, G., Goldwasser, M., Kao, M.-Y., Schweller, R.: Complexities for generalized models of self-assembly. In: Proceedings of the fifteenth annual ACM-SIAM Symposium on Discrete Algorithms, pp. 880-889 (2004)
- [4] Cheng, Q., de Espanes, P.M.: Resolving two open problems in the self-assembly of squares. Technical Report 793, University of Southern California (2003)
- [5] Kao, M.-Y., Ramachandran, V. DNA self-assembly for constructing 3D boxes. LNCS 2223, Proceedings of the 12th International Symposium on Algorithms and Computation (ISAAC 2001), P. Eades, T. Takaoka (Eds.), pp. 429-441 (2001)
- [6] Klavins, E.: Directed self-assembly using graph grammars. In: Proceedings of the Foundations of Nanoscience: Self Assembled Architectures and Devices, pp. 45-58 (2004)
- [7] Klavins, E.: Programmable self-assembly. IEEE Control Systems Magazine, 24(4), pp. 43-56 (2007)
- [8] Klavins, E., Ghrist, R., Lipsky, D. A grammatical approach to self-organizing robotic systems. IEEE Transactions on Automatic Control, 51(5), pp. 949-962 (2006)
- [9] Lathrop, J., Lutz, J., Patitz, M., Summers, S.: Computability and complexity in self-assembly. Theory of Computing Systems, to appear.
- [10] Lathrop, J., Lutz, J., Summers, S.: Strict self-assembly of discrete Sierpinski triangles. Strict self-assembly of discrete Sierpinski triangles, Theoretical Computer Science, 410 (2009), pp. 384-405.
- [11] Majumder, U., Reif, J., Sahu, S.: Stochastic Analysis of Reversible Self-Assembly. Journal of Theoretical and Computational Nanoscience 5(7), pp. 1289-1305, July 2008.
- [12] Naor, M., Stockmeyer, L.: What can be computed locally? SIAM Journal of Computing, 24(6), pp. 1259-1277 (1995)
- [13] Patitz, M., Summers, S.: Self-assembly of discrete self-similar fractals. To appear in: Proceedings of the Seventh International Conference on Unconventional Computation. Springer-Verlag, Heidelberg (2008)

- [14] Rothmund, P. W. K.: Theory and Experiments in Algorithmic Self-Assembly. Ph.D. thesis, University of Southern California, Los Angeles (2001)
- [15] Rothmund, P., Winfree, E.: The program-size complexity of self-assembled squares. In: Proceedings of the 32nd Annual ACM Symposium on Theory of Computing, pp. 459-468 (2000)
- [16] Seeman, N.: Denovo design of sequences for nucleic-acid structural-engineering. *Journal of Biomolecular Structure and Dynamics*, 8(3), pp. 573-581 (1990)
- [17] Soloveichik, D., Winfree, E.: Complexity of self-assembled shapes. *SIAM Journal of Computing*, 36(6), pp. 1544-1569 (2007)
- [18] Sterling, A. Distributed Agreement in Tile Self-Assembly. To appear in the Proceedings of the 15th International Meeting on DNA Computing and Molecular Programming, June 2009 (DNA 15), LNCS.
- [19] Sterling, A. Self-Assembling Systems are Distributed Systems. Tech report, July 2009. Submitted.
- [20] Wang, H.: Proving theorems by pattern recognition II. *Bell Systems Technical Journal*, 40, pp. 1-41 (1961)
- [21] Winfree, E.: Algorithmic Self-Assembly of DNA. Ph.D. thesis, California Institute of Technology, Pasadena (1998)